

CODED-APERTURE IMAGING IN NUCLEAR MEDICINE

Warren E. Smith
The Institute of Optics
University of Rochester

Harrison H. Barrett and John N. Aarsvold
Radiology Research Laboratory and
Optical Sciences Center
University of Arizona

SUMMARY

Coded-aperture imaging is a technique for imaging sources that emit high-energy radiation. This type of imaging involves shadow casting and not reflection or refraction. High-energy sources exist in x-ray and gamma-ray astronomy, nuclear reactor fuel-rod imaging, and nuclear medicine. Of these three areas nuclear medicine is perhaps the most challenging because of the limited amount of radiation available and because a three-dimensional source distribution is to be determined. In nuclear medicine a radioactive pharmaceutical is administered to a patient. The pharmaceutical is designed to be taken up by a particular organ of interest, and its distribution provides clinical information about the function of the organ, or the presence of lesions within the organ. This distribution is determined from spatial measurements of the radiation emitted by the radiopharmaceutical.

The principles of imaging radiopharmaceutical distributions with coded apertures will be reviewed. Included will be a discussion of linear shift-variant projection operators and the associated inverse problem. A system developed at the University of Arizona in Tucson consisting of small modular gamma-ray cameras fitted with coded apertures will be described.

INTRODUCTION

In nuclear medicine a radiopharmaceutical is given to a patient. The pharmaceutical is designed to go to a particular organ of interest, such as the brain, the heart, bone, or the liver, to name a few. The three-dimensional distribution of the pharmaceutical provides clinical information about how well the organ is functioning. This is quite different than the type of information provided by x-ray imaging (electron density), magnetic resonance imaging (MRI) (proton density and magnetization relaxation rates) and ultrasound (acoustic impedance of tissue). The distribution of the pharmaceutical is determined by imaging the radiation given off by the isotope that tags it. There is always a concern to limit the total amount of radiation that the patient is exposed to, so that in nuclear medicine we have a photon-limited situation.

The isotopes used in nuclear medicine fall into two broad categories: those that emit single gamma rays directly from the nucleus, and those that emit positrons from the nucleus. Three-dimensional imaging associated with the first category is called single photon emission computed tomography (SPECT), and is the subject of this paper. In this method the photons must be blocked by attenuating apertures. Imaging associated with the second category is called positron emission tomography (PET). In PET the positron that is emitted by a source nucleus annihilates an electron within 1 to 2 mm of the source point. This annihilation results in two photons, each of approximately 511 KeV, traveling in almost opposite directions. By coincidentally detecting these two photons with spatially separate detectors, the line along their path which contains the source point can be found. This technique removes the need to physically block the photons with apertures to determine their direction of origin. With PET one can obtain extremely good resolution by nuclear-medicine standards, on the order of 5 mm. The disadvantage of PET is that an on-site cyclotron is needed to create the short-lived positron-emitting isotopes. The expense associated with this requirement has limited the number of PET facilities. SPECT imaging, on the other hand, is relatively less expensive and well established throughout the world. Thus the motivation exists to continue to improve SPECT imaging techniques to approach the quality already attainable with PET.

Two-dimensional projections of source distributions are obtained in nuclear medicine by either scanning the source in two dimensions with a single, collimated gamma-ray point detector or by forming a two-dimensional image with a camera that is capable of measuring the x and y positions of the incident gamma rays and storing them in an image histogram. Such a camera is the Anger camera, named after its inventor (ref.1). This camera can also estimate the gamma-ray energy. Energy estimation is important for rejection of Compton-scattered radiation from the nuclear-medicine image. A photon that is Compton scattered by the attenuating tissues between the source and the detector will suffer an energy shift, dependent upon the angle of scatter. Fortunately in nuclear medicine the energy spectrum of the useful isotopes is relatively narrow, so that the Compton-scattered photons can be identified and removed if their energy is outside of the peak associated with the source isotope.

Gamma rays have such a high energy that they cannot be conveniently reflected or refracted. In front of the gamma-ray camera is thus placed a shadow-casting aperture, usually made of lead or some other high atomic-number element. There are two basic types of apertures, the collimator and the pinhole. The collimator consists of a large number of usually parallel holes drilled through a thick lead plate. Each hole causes the sensitivity of a given detector element to be confined to a narrow pencil that intersects the source distribution. This narrow pencil is an approximation to a line integral through the source. All of the holes together form a parallel-line 2-D projection of the source onto the 2-D detector. The pinhole is a single hole punched in a relatively thin lead plate. This aperture produces a pinhole image of the source distribution on the 2-D detector. The

pinhole image represents a series of line integrals through the object that converge on the pinhole. Conventional systems in nuclear medicine often employ parallel-hole collimators. The coded-aperture systems to be discussed employ arrays of pinholes.

Tomography in nuclear medicine is achieved by taking multiple views of the source distribution, and reconstructing a 3-D estimate of the source from these views. Conventionally these views are obtained by rotating a large gamma-ray camera fitted with a parallel-hole collimator around the patient. The camera stops every few degrees and takes a two-dimensional snapshot of the patient lasting about a minute. Each snapshot of the patient approximates a set of parallel line integrals, defined by the collimator, through the source volume at the particular angle. Neglecting attenuation of the source by the body, the set of all of these snapshots over 180 or 360 degrees constitutes an approximation to the Radon transform of the source distribution. The inverse Radon transform (ref.2) is then applied to these projections to form an estimate of the three-dimensional source distribution. This inverse involves filtering and then back-projecting each projection into the reconstruction space, and can be done rapidly with modern equipment. Without modification of the inverse Radon transform to include attenuation of the photons by the body, reconstructions appear darker for pixels deeper within the tomographic slice. This attenuation problem can be corrected analytically by the attenuated Radon transform (refs. 3,4), assuming constant attenuation and a known convex attenuation boundary. Typical scan times for the rotating-camera approach are 30 to 45 minutes. Dynamic studies of pharmaceutical uptake are ruled out because of the required motion of the camera about the patient.

In this paper we discuss tomography in nuclear medicine with non-moving coded apertures. The reconstruction of both 2-D and 3-D source distributions will be addressed. A coded-aperture system for nuclear medicine being developed at the University of Arizona will be described.

CODED-APERTURE TOMOGRAPHY IN NUCLEAR MEDICINE

In nuclear medicine we are able to observe only a small number of photons because the radiation dose to the patient is kept as low as possible and because the fractional solid angles of the collimator or pinhole openings are small, on the order of 10^{-5} . These openings must be small because in shadow-casting the ability of the aperture to resolve two closely spaced points in the source is directly proportional to the size of the openings. Thus we have a fundamental trade-off between the signal-to-noise ratio (SNR) in the nuclear-medicine image, which goes as the square root of the number of detected photons, and the resolution of the system. This trade-off is quite different in focusing systems, such as lenses focusing visible light, where the diffraction-limited spot size decreases (thus improving resolution) as the aperture is opened up, allowing more photons into the system.

There is thus strong motivation for increasing the number of photons in a nuclear-medicine image without degrading the resolution. To this end coded apertures have been developed. Figure 1 shows a single-view coded-aperture system. Here a planar source distribution is projected through an aperture consisting of several pinholes to form a coded image. The position of the pinholes represent the code. We have thus increased the number of photons detected by the system, at the price of overlap in the pinhole views of the object. This overlap is referred to as spatial "multiplexing", and is more serious for larger objects and denser spacing of the pinholes. Thus we suspect immediately that the code should be optimized with respect to the type of object that we wish to view.

In this planar case, neglecting radiometry and obliquity factors, we can write the coded image as a convolution of the source with the aperture:

$$g(x'', y'') \approx f(x''/m, y''/m) \otimes h(x''/M, y''/M) \quad (1)$$

where the double prime indicates detector coordinates. The quantity $g(x'', y'')$ is the coded image, $h(x''/M, y''/M)$ is the scaled aperture function, and $f(x''/m, y''/m)$ is the scaled source distribution. The source scaling $m = (z-d)/z$ and the aperture scaling $M = d/z$, where z is the source-aperture distance, and d is the source-detector distance. The two-dimensional convolution operator is represented by \otimes . As we see, both the source and the aperture functions are scaled differently in forming the coded image.

To form a reconstruction of the original source distribution, we use the concept of matched filtering. A matched filter is a version of the actual signal that we are looking for. It can be shown that a matched filter is the optimum filter to be used to detect a signal in the presence of noise (ref. 5). In the coded-imaging case, the matched filter is a properly scaled, inverted, complex-conjugated version of the original code, so that the reconstruction $\hat{f}(x'', y'')$ can be written as

$$\hat{f}(x''/m, y''/m) \approx g(x'', y'') \otimes h^*(-x''/M, -y''/M). \quad (2)$$

Equation (2) can be written, using Eq. (1), as:

$$\hat{f}(x''/m, y''/m) \approx f(x''/m, y''/m) \otimes [h(x''/M, y''/M) \otimes h^*(-x''/M, -y''/M)], \quad (3)$$

where the bracketed term represents the overall point-spread function (PSF) of the data-taking and reconstruction process, and is called the autocorrelation of the code. We must design the code to make its autocorrelation function as close to a delta function as possible, simultaneously allowing as many openings as possible. Unfortunately, these two requirements work against each other. Generally the autocorrelation has a large central peak surrounded by a background with structure that depends upon the number of openings. This background tends to both smear out the reconstruction as well as increase the noise in the reconstruction.

Note that we are not restricted to real and positive functions in our search for optimum codes. Any physically realizable code will be real and positive because of the shadow-casting nature of the image formation from the incoherent source. Bipolar complex codes can be simulated, however, by creating 4 separate codes and forming 4 separate coded images and suitably adding them in the computer with proper positive, negative, and imaginary weights. We pay the price for this flexibility by increasing the amount of time needed to form an image, however.

There has been considerable research into defining codes to optimize the SNR of the final reconstruction. Some of the more well-known codes are random pinhole arrays (ref. 6), the Fresnel zone plate (ref. 7), the annulus (ref. 8), and time-modulated apertures (ref. 9). Much of this code optimization has been in the context of single-view imaging of a planar object, however, as in Fig. 1. If we were to image a three-dimensional volume object with this approach, our reconstruction of Eq. (3) would be for a particular plane of the source, depending upon the scale factor used for the matched filter. The other planes of the source would present a strong background superimposed on this reconstruction, degrading both resolution and SNR of the plane of interest. Thus there is a fundamental limitation of the planar correlation decoding method described above because our basic data set consisting of a single view is not complete enough. We must in fact take multiple views of a volume object so that we are sampling its three-dimensional Fourier components sufficiently. Combining multiple views of the object to form a single volume reconstruction is not obvious with the planar decorrelation method described. In fact, we must generalize our entire approach to the problem and move away from the shift-invariant formulations of Eqs. (1-3).

With a multiple-view system, shown schematically in Fig. 2, we must give up the convenience of shift invariance. Thus the convolution operation can no longer be used to connect the object to the data. Instead the mapping from object to data takes on the more general form:

$$g(x'', y'', z'') = \int_{\text{source}} f(x, y, z) h(x'', y'', z''; x, y, z) d^3V, \quad (4)$$

where $g(x'', y'', z'')$ represents all of the coded images (spread out in three-dimensions), $f(x, y, z)$ is the three-dimensional source distribution, $h(x'', y'', z''; x, y, z)$ is the shift-variant mapping from the source to the coded images, and d^3V is a volume element of the source. All of the radiometry and aperture geometry is contained within $h(x'', y'', z''; x, y, z)$. The distribution $g(x'', y'', z'')$ forms the data set from which to reconstruct the estimate of the object $\hat{f}(x, y, z)$. Numerically, it is necessary to map the continuous problem into a discrete formulation by choosing a suitable basis set. We can see how this is done by a demonstration with a one-dimensional analog of Eq. (4):

$$g(x'') = \int_{\text{source}} f(x) h(x''; x) dx. \quad (5)$$

We can approximate $f(x)$ and $g(x'')$ each in the following way:

$$f(x) \approx \sum_{n=1}^N f_n \psi_n(x) \quad (6)$$

$$g(x'') \approx \sum_{m=1}^M g_m \phi_m(x'') \quad (7)$$

where

$$f_n \equiv \int_{-\infty}^{+\infty} f(x) \psi_n^*(x) dx \quad (8)$$

and

$$g_m \equiv \int_{-\infty}^{+\infty} g(x'') \phi_m^*(x'') dx''. \quad (9)$$

The basis sets $\psi_n(x)$ and $\phi_m(x'')$ span their respective spaces and are assumed orthonormal in this development. Thus we have approximated the source and the data with N and M discrete coefficients, respectively. An example of a particular source basis set is the "pixel" basis set, where the $\psi_n(x)$ are N non-overlapping shifted and scaled rectangle functions. Another example is the Fourier basis set, where the $\psi_n(x)$ represent complex exponentials, the eigenfunctions of shift-invariant operators. By the appropriate substitutions, we can now approximate Eq. (5) as:

$$g_m \approx \sum_{n=1}^N h_{mn} f_n \quad (10)$$

where

$$h_{mn} \equiv \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} h(x''; x) \psi_n(x) \phi_m^*(x'') dx dx''. \quad (11)$$

Thus the shift-variant problem of Eq. (5) is represented by a matrix multiplication, where the n^{th} column of the matrix \mathbf{H} , with elements

h_{mn} , represents the discrete, m -element shift-variant PSF due to the n^{th} expansion term of the source. In fact, if we choose the $\psi_n(x)$ and $\phi_m(x'')$ correctly, dependent upon $h(x''; x)$, we can have $h_{mn} = h_m$ for $m = n$, and $h_{mn} = 0$ otherwise. In other words, \mathbf{H} is diagonal or pseudo-diagonal if M is not equal to N . This choice of basis leads to what is called the singular value decomposition (SVD) of \mathbf{H} .

Generalizing to three dimensions, utilizing basis functions such as $\psi_n(x, y, z)$ and $\phi_m(x'', y'', z'')$, we can write Eq. (4) for all coded images in a way identical to Eq. (10). This can be done by simply ordering the N expansion coefficients of $f(x, y, z)$ into a one-dimensional $N \times 1$ vector \mathbf{f} and the M expansion coefficients of $g(x'', y'', z'')$ into a one-dimensional $M \times 1$ vector \mathbf{g} :

$$\mathbf{g} = \mathbf{H} \mathbf{f} + \mathbf{n}, \quad (12)$$

where we have introduced the $M \times 1$ zero-mean noise vector \mathbf{n} to allow for image degradation from effects outside of the direct mapping due to \mathbf{H} . Equation (12) is the general form of a shift-variant imaging system that we will use in the subsequent discussion of finding the source estimate $\hat{\mathbf{f}}$.

In general, Eq. (12) represents an ill-posed problem, in that one or more of the following conditions occur: no $\hat{\mathbf{f}}$ exists that satisfies \mathbf{g} exactly; $\hat{\mathbf{f}}$ is not unique; the solution $\hat{\mathbf{f}}$ is sensitive to small changes in \mathbf{g} or \mathbf{H} . We must usually content ourselves with a solution $\hat{\mathbf{f}}$ that agrees with \mathbf{g} to within some limits, and if these approximate solutions are not unique, choose one that satisfies some independent prior knowledge about \mathbf{f} . There are several techniques for finding $\hat{\mathbf{f}}$, such as singular value decomposition (SVD) alluded to briefly above (ref. 10), Monte Carlo methods (ref. 11), linear estimation theory (refs. 12, 13), and iterative methods (ref. 14). We will focus here on the Monte Carlo method, which we have found to be a practical technique for handling the large-scale pseudoinversion of Eq. (12) in the coded-aperture context. We have successfully simulated the reconstruction of volume objects \mathbf{f} of up to 32000 source elements from data sets \mathbf{g} consisting of nearly the same number of detector elements using less than 10 Mbytes of computer memory, in CPU times under 30 minutes on a VAX 8600. The reason for this space and time economy is that the \mathbf{H} matrix is sparse in coded-aperture imaging. Of course this sparseness is reduced as the number of pinhole openings increases, or as the size of the pinholes increases, since more detectors are being illuminated by each source element.

In the Monte Carlo reconstruction process we define an energy function E that is minimized when the reconstruction $\hat{\mathbf{f}}$ achieves a desired level of agreement with the data \mathbf{g} and simultaneously is consistent with any prior knowledge about the types of sources present. Such prior knowledge in the nuclear-medicine context consists of source positivity, source boundary, and perhaps correlation statistics between nearby source pixels. One of the cost functions that we have used is:

$$E = (1-\alpha) \parallel g - H \hat{f} \parallel^2 + (\alpha) \parallel \hat{f} - \langle \hat{f} \rangle \parallel^2 \quad (13)$$

where the double bar indicates magnitude of the vector and the brackets indicate an averaging process over nearest-neighbor pixels in the given estimate \hat{f} . The first term of this expression measures agreement with the data, and the second term imposes a smoothing constraint on \hat{f} , relating each pixel of the reconstruction to its nearest neighbors. The adjustable scalar α weights the agreement-with-data term against the smoothing term. We begin the reconstruction process with an initial guess at \hat{f} (a zero object or a uniform grey-level object). We then perturb each pixel of \hat{f} and calculate ΔE , the perturbation to E . This calculation is relatively rapid, because only a few detectors out of the hundreds or thousands of detectors actually see the perturbation to \hat{f} . It should be mentioned that only non-zero elements of H are required, so that even a 32000 by 32000 matrix can be stored in a small fraction of the space otherwise needed.

The perturbation is always accepted if $\Delta E \leq 0$, and if $\Delta E > 0$, it is accepted according to the Boltzmann probability of statistical mechanics:

$$P(\Delta E) = \exp(-\Delta E/kT) \quad (14)$$

where k is Boltzmann's constant (usually set to 1 in this context) and T is an effective "temperature" of the estimate at any given time. If T is large, we frequently allow large positive ΔE s into the reconstruction. If T is small, the probability of accepting large positive ΔE s is much reduced. The concept of starting the reconstruction at a large T and slowly reducing its value as E is decreased is known as "simulated annealing" (ref. 15). Such annealing is necessary if the energy surface E exhibits local minima: the occasional uphill energy swings of the reconstruction reduce the probability of being trapped in a local energy minimum. For quadratic energy functions as shown in Eq. (13), annealing is not required. However, if E is not quadratic, perhaps due to the imposition of strongly non-linear prior knowledge, annealing may become significant in improving the reconstruction. We have observed the importance of annealing for cases of very powerful prior knowledge, such as binary-object reconstruction, when a pixel is constrained to be on or off and the rules weighting its agreement with neighboring pixels are very nonlinear.

In our experience with the Monte Carlo algorithm, we find that we can typically obtain estimates \hat{f} that agree very well with the data, within a fraction of a percent. The smoothing constraint is a very important one; without it we get good data agreement, but there are large local fluctuations in the reconstruction that reduce its visual quality. The smoothing operation is imposed continuously as the reconstruction evolves permitting an ongoing compromise between the smoothing constraint and the data constraint.

An important aspect of coded-aperture imaging is the determination of the system operator \mathbf{H} . This matrix contains all of the geometry and radiometry (including attenuation, assuming known source-volume attenuation parameters) mapping the discrete object space to the discrete detector space. \mathbf{H} can be modeled theoretically as in Eq. (11), but for an actual system it should be found experimentally by a calibration procedure. Such a procedure consists of placing a point-source gamma-ray emitter in a volume attenuator that approximates the expected attenuating properties of the source, and stepping the point source through this volume one pixel location at a time. For each pixel location, the data set corresponds to a column of the \mathbf{H} matrix, including the effects of attenuation, radiometry, aperture vignetting, and detector efficiencies. It is important to have a bright enough source so that the SNR of the \mathbf{H} -matrix elements is high enough not to degrade the SNR of the reconstruction. Reconstructing the object using this \mathbf{H} matrix automatically includes the effects of attenuation and detector characteristics.

There are several advantages to pinhole coded-aperture imaging as compared to the conventional rotating collimated gamma-ray camera. The ability of a collimator to resolve two source points degrades faster with source depth than with a pinhole aperture. Thus the coded-images may contain higher spatial-frequency information than the collimator images. Also, the number of photons detected by a coded-aperture with many openings is greater than that of a collimator because the fractional solid angle of the coded aperture is greater. Thus we expect the SNR of a coded image to be superior to that of a collimator image. Finally, in a coded-aperture system consisting of multiple views, no detector motion is required so that dynamic studies are possible.

There are also disadvantages to the coded-aperture approach. Even though we detect more photons, this advantage is offset by the fact that we suffer from the multiplexing problem in the data sets. These two effects are coupled and both together determine the final SNR of the reconstruction. An additional complication is the need to carefully characterize the \mathbf{H} matrix through the calibration procedure described above. For a fixed system of modules and attenuation boundaries, however, this need be done only periodically. The attenuation boundaries can be fixed by placing the patient within a water sleeve, whose outer dimensions remain fixed. The reconstruction of the object from a coded-image data set is also more difficult in general than applying the inverse Radon transform in conventional tomography, but special-purpose hardware is being developed to optimize this procedure.

A set of small, independent gamma-ray cameras are being developed at the University of Arizona for applications in coded-aperture imaging (ref. 16). These cameras use a 10 cm by 10 cm NaI crystal coupled optically to 4 photomultiplier tubes (PMTs). The outputs of the 4 PMTs form a 20-bit address that extracts from a previously defined lookup table the statistically most likely x and y location of the gamma-ray impact point on the crystal face. Each camera, or a bank of cameras,

has its own coded aperture, thus forming a camera module. These modules can then be positioned about the patient in a configuration that will optimally utilize the detector area. Figure 3 is an example of an 8-view system for planar tomography that is currently being constructed in Arizona to be used for heart and brain imaging.

Preliminary simulations with systems similar to that of Fig. 3 demonstrate that state-of-the-art reconstructions are obtainable with data-acquisition times of the order of a third or less than that of the conventional rotating gamma-ray camera, which are typically 30 to 40 minutes. This potential data-acquisition time reduction, as well the static nature of the system allowing dynamic studies, may contribute to improving the state-of-the-art in nuclear-medicine imaging.

CONCLUSION

We have briefly described the principles of imaging in nuclear medicine, and have focused on a particular approach using coded apertures. The formulation of this shift-variant problem was developed, and a particular reconstruction algorithm was presented. A coded-aperture system being developed at the University of Arizona for tomographic imaging in nuclear medicine was briefly described.

ACKNOWLEDGMENTS

This work was supported by the National Cancer Institute through grant no. 2 P01 CA 23417. We thank Bruce Moore for his technical assistance.

REFERENCES

- 1) Anger, H.O., "Scintillation Camera," *Rev. Sci. Instrum.*, **29**, 27 (1958).
- 2) Barrett, H.H., and W. Swindell, *Radiological Imaging: The Theory of Image Formation, Detection, and Processing*, Vols. I and II (Academic, New York, 1981).
- 3) Tretiak, O., and C. Metz, "The Exponential Radon Transform," *SIAM. J. Appl. Math.* **39**, 341 (1980).
- 4) Clough, A.V., and H.H. Barrett, "Attenuated Radon and Abel Transforms," *J. Opt. Soc. Am. A*, **73**, 1590-1595 (1985).
- 5) Gaskill, J.D., *Linear Systems, Fourier Transforms, & Optics* (John Wiley, New York, 1978).
- 6) Dicke, R.H., "Scatter-hole Cameras for X-rays and Gamma Rays," *Astrophys. J.* **153**, L101 (1968).
- 7) Barrett, H.H., "Fresnel Zone Plate Imaging in Nuclear Medicine," *J. Nucl. Med.* **13**, 382-385 (1972).

- 8) Simpson, R.G., "Annular Coded-Aperture System for Nuclear Medicine," doctoral dissertation (University of Arizona, Tucson, Ariz., 1978).
- 9) Koral, K.F., W.L. Rogers, and F.G. Knoll, "Digital Tomographic Imaging with a Time-Modulated Pseudorandom Coded Aperture and an Anger Camera," *J. Nucl. Med.* 16, 402 (1975).
- 10) Strang, G. *Linear Algebra and Its Applications*, (Academic, New York, 1976).
- 11) Smith, W.E., R.G. Paxman, and H.H. Barrett, "Image Reconstruction from Coded Data: I. Reconstruction Algorithms and Experimental Results," *J. Opt. Soc. Am. A*, 2, 491-500 (1985).
- 12) Melsa, J.L., and D.L. Cohn, *Decision and Estimation Theory*, (McGraw-Hill, New York, 1978).
- 13) Smith, W.E., and H.H. Barrett, "Linear Estimation Theory Applied to the Evaluation of *A Priori* Information and System Optimization in Coded-Aperture Imaging," *J. Opt. Soc. Am. A*, 5, 315-330 (1988).
- 14) Frieden, B.R., "Image Enhancement and Restoration," in T.-S. Huang, ed., *Picture Processing and Digital Filtering*, Vol. 6 of Topics in Applied Physics, (Springer-Verlag, New York, 1975).
- 15) Kirkpatrick, S., C.D. Gelatt, Jr., and M.P. Vecchi, "Optimization by Simulated Annealing," *Science*, 220, 671-680 (1983).
- 16) Aarsvold, J.N., H.H. Barrett, J. Chen, A.L. Landesman, T.D. Milster, D.D. Patton, T.J. Roney, R.K. Rowe, R.H. Seacat, III, and L.M. Strimbu, "Modular Scintillation Cameras: A Progress Report," *Medical Imaging II: Image Formation, Detection, Processing, and Interpretation*, SPIE 914, 319-325 (1988).

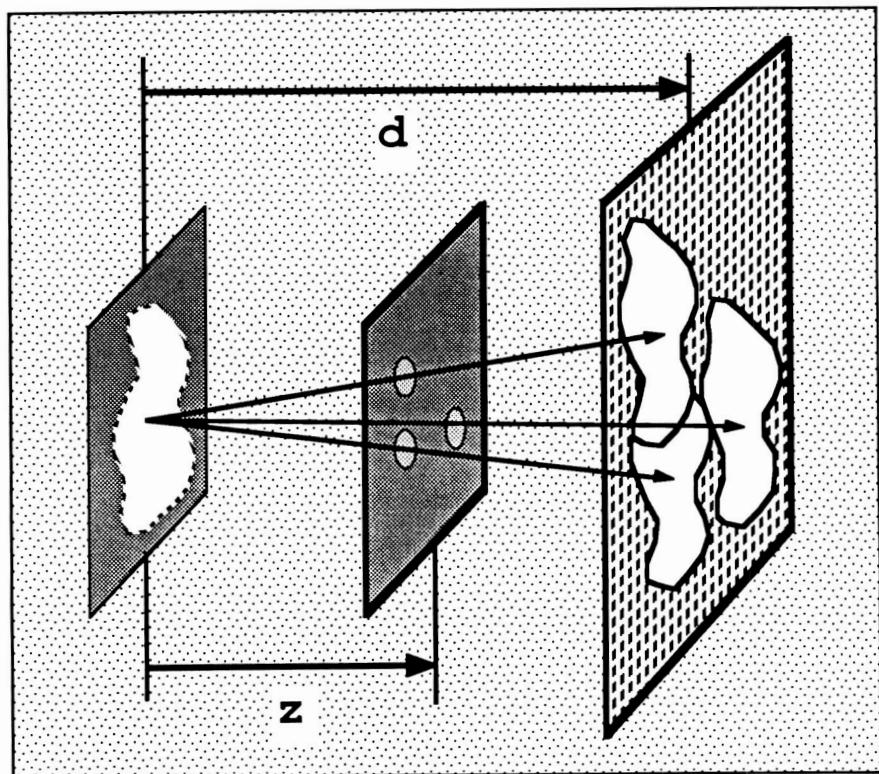


Figure 1) A single-view coded-aperture system,
imaging a planar source.

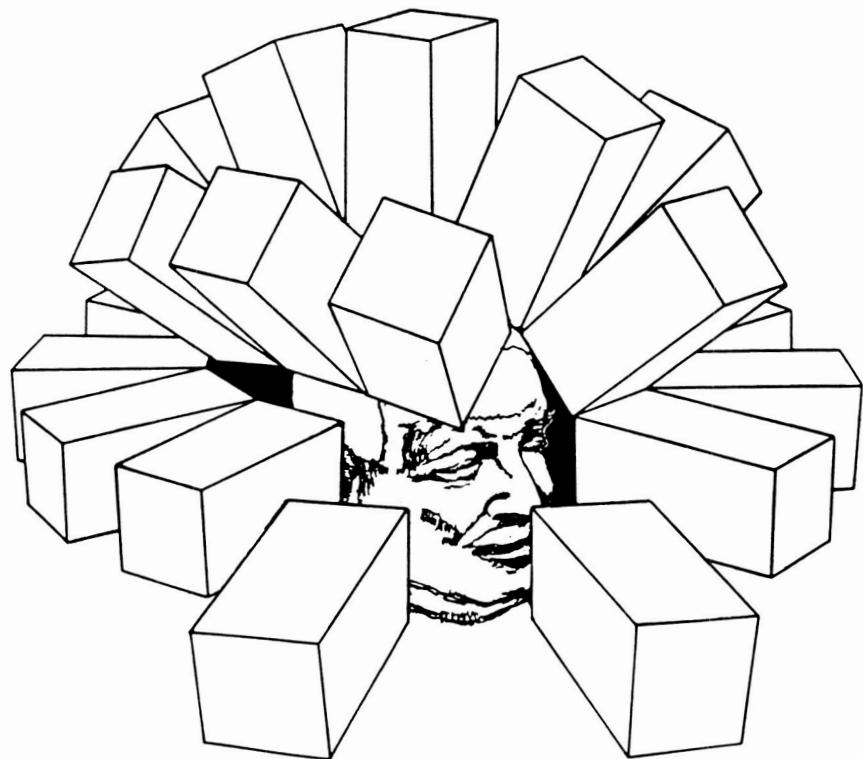


Figure 2) A multiple-view coded-aperture system,
imaging a volume source.

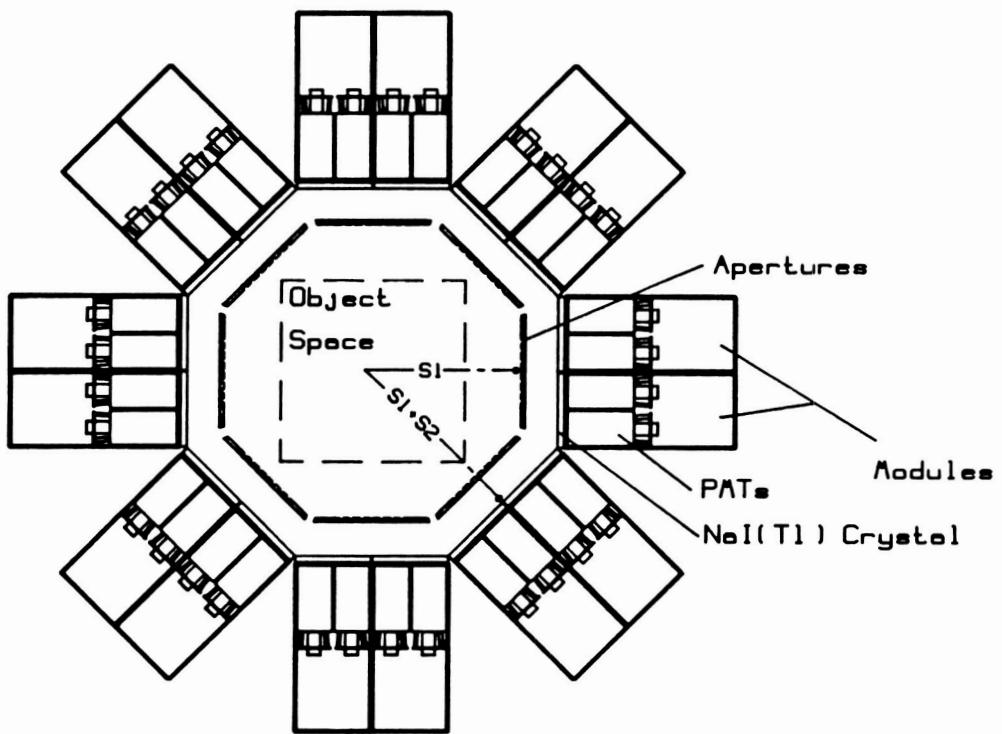


Figure 3) An octagonal coded-aperture system being used at the University of Arizona for planar tomography.